

METHOD AND APPARATUS FOR IMPROVING IMAGE APPEARANCE

BACKGROUND OF THE INVENTION1. Field of Invention

5 [001] This invention relates to systems and methods for improving the appearance of captured images.

1. Description of Related Art

10 [002] In the digital reproduction of documents, a bitmap is created which may be described as an electronic image with discrete signals, i.e. pixels, defined by a position and a density. In conventional image capture devices, such as facsimile and scanner devices, image degradation of captured bilevel image data often occurs. This degradation, such as lower resolution, noise, change in contrast and the like, is well within the visual acuity of the human eye. If the captured image data is output to a recording medium without adjusting for the degradation, the outputted image will include the degradation. Even though such bilevel images are usually readable, they are often difficult or unpleasant to read. Such images are also not presentable for formal purposes. This is because the human eye can sense this image degradation, and the perceived quality of the resulting image suffers greatly even for small degradation.

15 [003] Various attempts at remedying such problems have been performed. An example is U.S. Patent No. 5,303,313 to Mark et al., which provides a method of image enhancement through use of a compressed representative image. Another example is described in J.D. Hobby et al., "Enhancing degraded document images via bitmap clustering and averaging," *ICDAR '97: Fourth Int. Conference on Document Analysis and Recognition, 1997*. Both Patent No. 5,303,313 and the Hobby article provides a basic strategy. In Hobby, the strategy includes: clustering bitmaps, computing representatives for each cluster, ~~reclustering~~, and then assembling an output. For initial clustering, Hobby uses a feature-based approach. To computer cluster representatives, Hobby uses a method that aligns the scans by centroids of black pixels, sums the scans to give a histogram, smooths the histogram to give a gray-level representative, and determines a polygonal outline that stays within a certain gray "tube" yet has a minimum number of inflection points. This computation method is described in J.D. Hobby and H.S. Baird, "Degraded Character Image

20
25
30

Restoration", *Proc. 5th Annual Symp. On Document Analysis and Image Retrieval*, 1996, pps. 177-189. To align and form the assembled output, Hobby appears to use the alignment computed when computing cluster representatives. Patent No. 5,303,313 does not perform any reclustering, and instead is concerned primarily with compression.

[004] While the Hobby method shows some improvement in images and increases resolution, there are many refinements that can be made.

SUMMARY OF THE INVENTION

[005] Methods and systems of this invention improve the appearance of a captured bilevel image to enable better reading and improved downstream processing, such as deskewing or optical character recognition (OCR).

[006] The methods and systems of this invention separately avoid image degradation that appear in the captured bilevel image.

[007] This invention separately provides systems and methods for printing images that reduce image degradation introduced during image capturing to provide a printed image with improved appearance.

[008] This invention separately provides systems and methods that have more reliable initial clustering, a reduction of clusters without introducing new errors, super-resolved placement of representatives, and other image enhancement including breaking-up of run-together letters of text.

[009] In various exemplary embodiments of the methods and systems according to this invention, the output image may have higher resolution of the input image.

[0010] In various exemplary embodiments of the methods and systems according to this invention, a bitmap representation of a captured image is clustered into a plurality of clusters, representatives of the clusters are determined, the bitmap may then be reclustered, and then an output image is assembled.

[0011] In various exemplary embodiments of the methods and system according to this invention, connected components of dark pixels are clustered from across the image, and a "most likely" representative image for each cluster of images is determined, with likelihood determined by a probabilistic model of the image capturing process. The representative images are themselves bitmaps. In various exemplary embodiments, the representative images are at higher resolution. These

representative images may be reclustered and finally assembled in an output page by replacing each member of a cluster by the cluster's representative.

[0012] In various exemplary embodiments of the methods and systems according to this invention, initial clustering uses a Hausdorff matching algorithm.

5 [0013] In various exemplary embodiments of the methods and systems according to this invention, cluster representations are determined by using a hill-climbing optimization procedure to approximate the most probable double-resolution representative. This has the advantage that it can rigorously incorporate Bayesian priors and learned or guessed scanner distortion parameters resulting in more accurate
10 sharp features and reliable overall blackness. However, other optimization procedures can be substituted.

[0014] In various exemplary embodiments of the methods and systems according to this invention, reclustered combines or eliminates clusters but does not split clusters, thus reducing the total number of clusters.

15 [0015] In various exemplary embodiments of the methods and systems according to this invention, the assembly replaces representatives in their likeliest positions.

[0016] In various exemplary embodiments of the methods and system according to this invention, *a priori* (prior) probability distributions on bitmaps are
20 used to determine the most likely representative images.

[0017] In various exemplary embodiments of the methods and system according to this invention, *a priori* probability distributions based on so-called chain codes may be implemented. The methods and systems of this invention then use the representative images to recluster connected components, and finally to assemble the
25 output page by replacing each member of a cluster of images by that cluster's representative image. Thus, the degradation is reduced or eliminated, and an improved bilevel image is obtained.

[0018] In various exemplary embodiments of the methods and systems of the invention, improved deskewing and optical character recognition (OCR) with
30 improved accuracy can be attained.

[0019] These and other features and advantages of this invention are described in or are apparent from the following detailed description of various exemplary embodiments.

10046195-01602

BRIEF DESCRIPTION OF THE DRAWINGS

[0020] Various exemplary embodiments of this invention will be described in detail, with reference to the following figures, in which:

[0021] Fig. 1 shows one exemplary embodiment of a system that includes an image processing apparatus and the image capture device according to this invention;

[0022] Fig. 2 shows one exemplary embodiment of the image improvement circuit or routine of Fig. 1;

[0023] Fig. 3 shows one exemplary embodiment of the image capture device of Fig. 1;

[0024] Fig. 4 illustrates an example of the default point spread function of a sensor shown in Fig. 3;

[0025] Fig. 5 illustrates an example of the probability curve as a sigmoidal function of the weight of input black of the sensor shown in Fig. 3;

[0026] Fig. 6 is a flowchart outlining one exemplary embodiment of a method for processing an image according to this invention;

[0027] Fig. 7 is a flowchart outlining one exemplary embodiment of the image improvement data determining step of Fig. 6;

[0028] Fig. 8 shows a comparison of text from a fine 200 dpi fax input, image enhancement using a known method, and image enhancement according to the invention;

[0029] Fig. 9 shows details of a flatbed scanner input, the input lightened, and the input darkened;

[0030] Fig. 10 shows details of a fax input, the fax input after 1 round of pixel flipping according to the invention, and after 4 rounds of pixel flipping according to the invention; and

[0031] Fig. 11 shows details of the fax input of Fig. 10 after 4 rounds of pixel flipping using priors, after reclustering, and after breaking up run-together letters.

DETAILED DESCRIPTION OF PREFERRED EMBODIMENTS

[0032] Various exemplary embodiments of the invention will be described, each of which can provide image improvement to captured images. In these various embodiments, connected components from across one or more pages of a captured bilevel image are clustered and a "most likely" representative for each cluster is computed by a probabilistic model of the scanning process. Representative images

are themselves bitmaps, but may be at a higher resolution. These representations are then used to re-cluster connected components. An output image is then assembled by replacing each family member of a cluster by the cluster's representative.

[0033] The invention may be implemented on the exemplary system shown in Fig. 1. As shown in Fig. 1, an image capture device 100 and an input device 120 are connected to an image processing apparatus 200 over links 110 and 122, respectively. Similarly, an image data sink 300 can be connected to the image processing apparatus 200 over a link 310.

[0034] The image capture device 100 can be a digital camera, a scanner, a facsimile machine, a digital copier, or any other known or later developed device that is capable of capturing an image and generating electronic image data that has been captured according to the image capture techniques described above. Similarly, the image capture device 100 can be any suitable device that stores and/or transmits electronic image data such as a client or a server of a network that has been captured according to the image capture techniques described above.

[0035] The image capture device 100 can be integrated with the image processing apparatus 200, as in a digital copier or a facsimile machine having an integrated scanner. Alternatively, the image capture device 100 can be connected to the image processing apparatus 200 over a connection device, such as a modem, a local area network, a wide area network, an intranet, the Internet, any other distributed processing network, or any other known or later developed connection device.

[0036] It should also be appreciated that, while the electronic image data can be generated at the time of printing an image from electronic image data, the electronic image data could have been generated at any time in the past. The image capture device 100 is thus any known or later developed device that is capable of supplying electronic image data that has been captured according to the image capture techniques described above over the link 110 to the image processing apparatus 200. The link 110 can thus be any known or later developed system or device for transmitting the electronic image data from the image capture device 100 to the image processing apparatus 200. Non-limiting examples include a direct cable connection, a connection over a wide area network or a local area network, a connection over an intranet, a connection over the Internet, or a connection over any other distributed processing network or system. In general, the link 110 can be any known or later

developed connection system or structure usable for connection between two components to transmit data.

[0037] The input device 120 can be any known or later developed device for providing control information from a user to the image processing apparatus 200.

Thus, the input device 120 can be a control panel of the image processing apparatus 200, or a control program executing on a locally or remotely located general purpose computer or the like. As with the link 110 described above, link 122 can be any known or later developed device for transmitting control signals and data input using the input device 120 from the input device 120 to the image processing apparatus 200.

[0038] The image data sink 300 can be any known or later developed device that can receive the reconstructed composite image from the image processing apparatus 200. Thus, the image data sink 300 can be a display, an image data sink such as a laser printer, a digital copier, an inkjet printer, a dot matrix printer, a dye sublimation printer, or the like. The image data sink 300 can also be any known or later developed storage device, such as a floppy disk and drive, a hard disk and drive, a writeable CD-ROM or DVD disk and drive, flash memory, or the like. It should also be appreciated that the image data sink 300 can be located locally to the image processing apparatus 200 or can be located remotely from the image processing apparatus 200. Thus, like the links 110 and 122, link 310 can be any known or later developed connection system or structure usable to connect the image processing apparatus 200 to the image data sink 300. Specifically, the link 310 can be implemented using any of the devices or systems described above with respect to links 110 and 122.

[0039] In general, the image data sink 300 can be any known or later developed device that is capable of receiving data output by the image processing apparatus 200 and either storing, transmitting or displaying the data. Thus, the image data sink 300 can be either or both of a channel device for transmitting the data for printing, display or storage or a storage device for indefinitely storing the data until there arises a need to print, display or further transmit the data.

[0040] If data sink 300 is a channel device, it can be any known structure or apparatus for transmitting data from the image processing apparatus 200 to a physically remote storage or display device. Thus, the channel device can be a public

switched telephone network, a local or wide area network, an intranet, the Internet, a wireless transmission channel, any other distributing network, or the like. Similarly, the storage device can be any known structural apparatus for indefinitely storing image data such as a RAM, a hard drive and disk, a floppy drive and disk, an optical drive and disk, a flash memory or the like. For example, the image data sink 300 may be a printer, a facsimile machine, a digital copier, a display, a host computer, a remotely located computer, or the like.

[0041] As shown in Fig. 1, the image processing apparatus 200 includes a controller 210, an input/output interface 220, a memory 230, an image improvement circuit or routine 240 and an image processing circuit or routine 250, each of which is interconnected by a control and/or data bus 260. The links 110, 122 and 310 from the image capture device 100, the input device 120, and the image data sink 300, respectively, are connected to the input/output interface 220. The electronic image data from the image capture device 100 and any control and/or data signals from the input device 120 are input through the input interface, and, under control of the controller 210, are stored in the memory 230.

[0042] The memory 230 preferably has at least an alterable portion and may include a fixed portion. The alterable portion of the memory 230 can be implemented using static or dynamic RAM, a floppy disk and disk drive, a hard drive, flash memory, or any other known or later developed alterable volatile or non-volatile memory device. If the memory includes a fixed portion, the fixed portion can be implemented using a ROM, a PROM, an EPROM, and EEPROM, a CD-ROM and disk drive, a writable optical disk and disk drive, or any other known or later developed fixed memory device.

[0043] The image improvement circuit 240 inputs signals received from the image capture device 100. The image improvement circuit 240 then outputs image improvement data, which can have a higher resolution than the originally received image data corresponding to the original document, to the image processing circuit or routine 250. The image processing circuit or routine 250 adjusts the captured image data to generate improved image data from the originally received image data, based on the image improvement data from the image improvement circuit or routine 240.

[0044] The processed image data is outputted from the image processing apparatus 200 to the image data sink 300 over the link 310. The image processing

circuit 250 can also process the improved image data to apply any other known or later developed image processing technique. Accordingly, when the improved image data is output to the image data sink 300, the resulting image can contain any additional known or later developed image enhancements.

5 **[0045]** The image processing apparatus 200 shown in Fig. 1 is connected to the image data sink 300 over the link 310. Alternatively, image data sink 300 may be an image output terminal that is integral part of the image processing apparatus 200. An example of this configuration would be a digital copier or the like. It should be appreciated that the image processing apparatus 200 can be any known or later
10 developed type of image processing apparatus. There is no restriction on the form the image processing apparatus 200 can take.

[0046] As indicated above, the image data sink 300 may be an integrated device with the image processing apparatus 200, such as a digital copier, computer with a built-in printer, or any other integrated device that is capable of producing a
15 hard copy image output. However, as another example, the image processing apparatus 200 and the image data sink 300 may be physically separate, such as a computer memory and a printer.

[0047] After being processed by the image processing apparatus 200, the image data is output to the image data sink 300. The data may be stored in the
20 memory before, during and/or after processing by the image processing apparatus 200, as necessary.

[0048] It should be understood that various components of the image processing apparatus 200 shown in Fig. 1, such as the image improvement circuit or routine 240, the image processing circuit or routine 250, and the controller 210, can
25 each be implemented as software executed on a suitably programmed general purpose computer, a special purpose computer, a microprocessor or the like. In this case, these components can be implemented as one or more routines embedded in a printer driver, as resources residing on a server, or the like. Alternatively, these components can be implemented as physically distinct hardware circuits within an ASIC, or using an
30 FPGA, a PDL, a PLA, or a PAL, or using discrete logic elements or discrete circuit elements. The particular form each of the components shown in Fig. 1 will take is a design choice and will be obvious and predictable to those skilled in the art.

[0049] In one exemplary embodiment of this invention, the image improvement circuit or routine 240 is able to initially cluster portions of the bitmap of the received image data. From this data, the image improvement circuit or routine 240 determines the representative images for each of the clusters. Initial clustering is preferably attained using a Hausdorff matching method. A suitable example of such can be found in U.S. Patent No. 5,835,638 to Rucklidge et al., the disclosure of which is incorporated herein by reference in its entirety. See also the DigiPaper article at <http://www.cs.cornell.edu/digipaper>. Other methods of initial clustering are known and could be substituted. An exemplary other known method of determining the initial clustering is described in J.D. Hobby et al., "Enhancing degraded document images via bitmap clustering and averaging," *ICDAR '97: Fourth Int. Conference on Document Analysis and Recognition, 1997*. However, this latter method may be less reliable.

[0050] Fig. 2 shows one exemplary embodiment of the image improvement circuit or routine of this invention. As shown in Fig. 2, in the image improvement circuit or routine 240, image data is input to a bitmap clustering portion 242. In the bitmap clustering portion 242, portions of the bitmap of the received image data are initially clustered and the clustered data is input to a representative determining portion 244. From this data, the representative determining portion 244 determines representative images for each of the clusters. In a bitmap reclustering portion 246, the bitmap is then reclustered. Then, the bitmap is reassembled by replacing each member of a cluster by the representative for that cluster.

[0051] The reclustered bitmap is output to the image processing circuit or routine 250. The image processing circuit or routine 250 then assembles an improved version of a captured document from this determination with a higher resolution or improved appearance.

[0052] Fig. 3 shows one exemplary embodiment of the image capture device of this invention. As shown in Fig. 3, the image capture device 100 includes a rectangular grid of point sensors 102 that sample an original image. The outputs of the sensors 102 are black and white (or dark and light) pixels.

[0053] In this exemplary embodiment, each sensor 102 detects a roughly disk-shaped region 104 of the original image and outputs a white or black pixel based on the sensed image density of the detected region of the scanned document. In

general, the likelihood that a particular sensor 102 will output a black pixel is probabilistically dependent upon the total weight of black in the detected region. Although each sensor 102 is preferably positioned at the center of the pixel, it should be appreciated that the sensors 102 may be positioned at the corners rather than at the centers of the input pixels, and hence a point spread function can have four center coefficients.

[0054] The coefficients for the pixels within the disk-shaped region define a point spread function of the sensors 102. A curve showing the probability that an output pixel of a sensor 102 is black defines the response function of that sensor 102. Fig. 4 shows one exemplary point spread function of a sensor 102. Fig. 5 shows one exemplary probability curve of a sensor 102. In this example, the probability curve is a sigmoidal function of the weight of input black.

[0055] The response function can be varied to model different threshold settings for the image capturing device 100. As shown in Fig. 5, a sigmoid symmetric around 0.5 implies no gain for the image capturing device. That is, the expected amount of output black equals the amount of input black. As shown in Fig. 5, a sharp sigmoid upward slope between .2 and .6 models an image capturing device with some gain.

[0056] If optical characteristics of the sensors 102 are known or can be inferred, then the point spread and response functions can be set specifically for a given input image. Alternatively, a user can control the point spread and response functions using the input device 120.

[0057] The bitmap clustering portion 242 of the image improvement circuit or routine 240 initially clusters the portions of the received bitmap image into a plurality of clusters of portions, using Hausdorff matching. The representative determining portion 244 determines the representative images for the clusters. The bitmap recluster portion 246 then reclusters the received bitmap image using the scanner model which contains the point spread function and probability curve of sensors 102. The image processing circuit or routine 250 assembles the output image using the reclustered bitmap from the image improvement circuit or routine 240.

[0058] In particular, the "portions" are connected components of black pixels, and are clustered from across the image as discussed above, preferably using a Hausdorff matching algorithm. A connected component is an island of dark (black in

the case of a binary black/white image) pixels in a binary scan of a document. That is, a set of dark pixels connected diagonally or orthogonally and surrounded by white. A "most likely" representative image for each cluster of portions is then determined. In various exemplary embodiments, the likelihood of the "most likely" representative image is determined by a probabilistic model of the image capturing process. An approximate most likely representative can be discovered by a hill-climbing optimization procedure. The representative images are themselves bitmaps, with representative bitmaps being at higher resolution than the underlying received image data.

[0059] In one exemplary embodiment of the methods and systems according to this invention, the representative determining portion 244 of the image improvement circuit or routine 240 uses an *a priori* (prior) probability distribution on the bitmap portions to determine the most likely representative image of each cluster of portions. The *a priori* probability distribution is based on "chain codes". A "chain code" is a sequence of North, South, East and West directions taken while traversing the boundary of a connected component. For more information on chain codes, see co-pending Appl. No. 09/749,690, filed December 28, 2000, the subject matter of which is incorporated herein in its entirety.

[0060] The bitmap reclustering portion 246 of the image improvement circuit or routine 240 then uses the representative images to recluster the portions, such as the connected components. The image processing circuit or routine 250 assembles the output page by replacing each member of a cluster by that cluster's representative. Thus, for example, image degradation is reduced or eliminated, and an improved bilevel image may be obtained.

[0061] Fig. 6 is a flowchart outlining one exemplary embodiment of an image processing method according to this invention. Beginning at step S1000, control advances to step S1100, where the document is input. Then, in step S1200, an image of the document is captured. Next, in step S1300, image improvement data is determined based on the captured image. Control then advances to step S1400. In step S1400, the captured image data is adjusted to provide the improved image data. Next, in step S1500, the adjusted image data image is output as output data. Then, in step S1600, the process stops.

[0062] Fig. 7 is a flowchart outlining one exemplary embodiment of the image improvement data determination step S1300. Beginning in step S1300, control advances to step S1310, where adjacent pixels are analyzed to determine connected components and to define each connected component as a cluster. Then, in step S1320, the clusters are pair-wise compared to determine if a match is found. If a match is found, control continues to step S1330. Otherwise, the clusters do not match, and control jumps to step S1340.

[0063] In step S1330, the matched clusters are combined into a corresponding cluster. Next, in step S1340, it is determined whether all clusters have been analyzed. If not, flow returns to step S1320. If all have been analyzed, flow advances to step S1350 where a representative image for each cluster is found. Then, in step S1360, reclustering and image reassembling is performed by replacing each of the members of a cluster with that cluster's representative image. Control then advances to step S1370, where control returns to step S1400.

[0064] A more detailed explanation of the invention will now be described. As described previously, the inventive methods and systems for performing image improvement include the process steps of: 1) initial clustering; 2) finding representatives; 3) reclustering; and 4) assembling the output. Each will be described in detail below.

[0065] As shown in Fig. 1, and previously discussed, the image processing apparatus 200 is preferably implemented on a programmed general purpose computer. However, the image processing apparatus 200 can also be implemented on a special purpose computer, a programmed microprocessor or microcontroller and peripheral integrated circuit elements, an ASIC or other integrated circuit, a digital signal processor, a hardwired electronic or logic circuit such as a discrete element circuit, a programmable logic device such as a PLD, PLA, FPGA or PAL, or the like. In general, any device, capable of implementing a finite state machine that is in turn capable of implementing the four basic process steps or the flowcharts shown in Figs. 6 and 7, can be used to implement the image processing apparatus 200.

[0066] A. Initial Clustering:

[0067] When the portions of the received image are implemented as connected components, initial clustering is performed by connecting each black pixel in the captured image to an adjacent black pixel. A connected component is a

maximal set of black pixels in the initial binary raster, such that each black pixel is connected to each other by a path of adjacent black pixels. In various exemplary embodiments, the adjacent pixel can include a diagonally adjacent pixel. As such, each pixel may have a total of 8 neighbors.

5 **[0068]** In one exemplary embodiment, matching is used to form family members for the clusters. In matching, initially, each connected component is in a cluster of its own and thus is that cluster's representative image. Clusters are then combined by finding matching representative images. As the cluster membership changes, either by combining clusters or by dropping members that no longer match
10 the representative image, cluster representative images are redetermined by thresholding aligned histograms. In various exemplary embodiments, the threshold can be set to preserve median blackness.

[0069] For any two connected components A and B, a bounding box is just formed around each connected component A and B. Then, the connected components
15 A and B are aligned to each other by aligning the centers of their bounding boxes. The connected components A and B will match each other if:

$$|A| - |A \cap \overline{B}| \leq f(|\partial A|) \text{ and } |B| - |B \cap \overline{A}| \leq f(|\partial B|)$$

where:

 |A| denotes the number of black pixels in A;
20 $A \cap B$ denotes the pixels that are black in both A and B;
 \overline{A} denotes a one-pixel dilation of the black pixels in A;
 ∂A denotes the boundary of A, that is, the set of black pixels with white neighbors; and
 $f(n)$ equals 0 for $n \leq 3$, and $.025n$ for $n \geq 7$, and interpolates between these two
25 lines for $3 < n < 7$.

[0070] A dilation of the black pixels in A is the component that has a black pixel wherever A has either a black pixel or a white pixel orthogonally bordering a black pixel. For example, a topology-preserving dilation is used, which refuses to blacken a pixel if it would join two connected components in its 8-neighborhood.

30 **[0071]** In other words, for A and B to match, the number of pixels of A lying outside B must be very small, and vice versa. In various exemplary embodiments, an

additional test can be used to stop a match if either $A \setminus \overline{B}$ or $B \setminus \overline{A}$ includes a set of more than three black pixels that can be enclosed by a 3 x 3 box.

[0072] Generally, this initial clustering using the Hausdorff method uses a distance measuring technique that is a measure for comparing point sets that can be used to compare binary images. Further details of initial clustering using this Hausdorff matching method can be found in U.S. Patent No. 5,835,638 to Rucklidge et al. and the DigiPaper article (<http://www.cs.cornell.edu/digipaper>), the disclosures of which are incorporated herein by reference in their entirety.

[0073] B. Finding Representatives:

[0074] Using the connected components, optimal representative images of the clusters are determined. Briefly, this is determined by a hill-climbing approach. However, before optimal representatives can be better explained, it first must be explained how to compute the probability that a given connected component A is a scan of a given original image B. τ represents a translation of the scanner's sensor grid with respect to B. $w_{ij}(\tau)$ denotes the weight of black in B as seen by the sensor in row i and column j . Using the exemplary point spread and response functions given in Fig. 4 (or other known point spread for a particular sensor), the probability $p(w_{ij}(\tau))$ that the sensor's output pixel will be black can be computed. The probability that the pixel in row i and column j has value A_{ij} (black or white) given B and τ is determined as:

$$P[A_{ij}|B, \tau] = \begin{cases} p(w_{ij}(\tau)) & \text{if } A_{ij} \text{ is black;} \\ 1 - p(w_{ij}(\tau)) & \text{if } A_{ij} \text{ is white.} \end{cases} \quad (1)$$

where:

τ represents a translation of the sensor grid with respect to the given original image region B;

$w_{ij}(\tau)$ denotes the weight of black in the given original image region B seen by the sensor in row i and column j ; and

$p(w_{ij}(\tau))$ denotes the determined probability that the sensor's output pixel would be black.

[0075] In one exemplary embodiment, the sensors 102 act independently. That is, the randomization of the response function is independent from sensor to

sensor. Thus, the individual pixel probabilities can be multiplied to give the probability $P[A | B, \tau]$ that the connected component A is a capture of the given original image region B at translation τ as:

$$P[A|B, \tau] = \prod_{ij} P[A_{ij}|B, \tau] \quad (2)$$

[0076] The connected component A and the given original image region B are each padded with white pixels, and indices i and j run over all positions in the union of the bounding boxes of the connected component A and the given original image region B.

[0077] The above equations (1) and (2) assume a specific translation τ . However, since τ is unknown, τ can be optimized over all possible translations as:

$$P[A|B] = \max_{\tau} P[A|B, \tau]. \quad (3)$$

[0078] While this may involve a difficult optimization problem, if the connected component A and the given original image region B have been pre-aligned by the centroids of their bounding boxes, then the determination may be limited to the nine shortest vectors in this lattice. That is, the determination may be limited to a shift of -1, 0, or 1 in each of the x - and y -coordinates.

[0079] The probability of an entire cluster of bitmaps C is determined by multiplying the probabilities of each individual bitmap. Since probabilities become very small, logarithms are added as:

$$P[C|B] = \prod_{A \in C} P[A|B] \quad (4)$$

[0080] The optimal representative image of the given original image region B for a cluster C is the one that maximizes $P[C|B]$. This is represented as:

$$P[B|C] = P[C|B] \cdot \frac{P[B]}{P[C]} \quad (5)$$

[0081] Probability $P[B]$ is the *a priori* probability of the image of the given representative original image region B , which is assumed to be the same for all given original image regions B , and $P[C]$ is the *a priori* probability of the cluster C , which is constant.

5 [0082] To find the B that maximizes $P[C|B]$, a hill-climbing approach is preferably used. The initial representative image B^0 of the original image region B is simply the cluster representative image with each pixel split into four double-resolution pixels. For each captured connected component A in the cluster, the translation τ for $P[A|B^0, \tau]$ is determined by searching the 9 shortest vectors as above. Next, $P[C|B^0]$ is determined and recorded. The translated image capture is summed to form a double-resolution histogram, which is used to guide the search for the representative. Pixels in the initial representative image B^0 are flipped. That is, pixels in the initial representative image B^0 are changed from white to black or vice versa, based on this histogram.

10 [0083] To determine the next representative image B^1 , only the most clearly indicated flips are used. Specifically, in various exemplary embodiments, only these white pixels are flipped where, for example, more than 60% of the captured connected component have black pixels at the corresponding location and these black pixels are flipped where fewer than 40% of the captured connected component have white pixels at the corresponding location. Image capture with respect to B^1 is then aligned, $P[C|B^1]$ is determined and recorded, and the histogram is updated.

15 [0084] For the next representative image B^2 , flipping is a little more aggressive, with white pixels over 55%, for example, and black pixels under 45% being flipped, and the alignment and updating cycle are repeated. For the next representative image B^3 and subsequent representatives, pixels are flipped according to the expected number in the corresponding histogram bin, rather than by fixed percentages. If the observed number exceeds the number predicted by the scanner model by more than a certain percentage or number, a white pixel is flipped to black, and vice versa. The process is halted either when no pixels flip or after a fixed number of cycles. In various exemplary embodiments, the process is halted after four cycles. The representative image is the B^i with maximum $P[C|B^i]$.

20 [0085] This ad hoc optimization heuristic starts out with conservative flips and then gradually becomes more aggressive. This is because flipping a pixel tends to

inhibit its neighbors from flipping. Hence only the "locally most flippable" pixels are used to qualify in the early rounds. On the other hand, a more sequential approach, such as flipping pixels one at a time starting from the "most flippable" would be unacceptably slow. As a further way to speed up the process, new alignments are not
 5 determined after the representative image B^2 or determined subsequently. Typically, there is a lot of flipping from the representative image B^0 to the representative image B^1 , and only a little bit of fine tuning-which rarely changes the alignments-in subsequent rounds.

[0086] The cluster representative for a large cluster (at least five scans)
 10 typically has a noticeably better appearance than a representative for a small cluster. Fig. 10 shows some examples of this. The left section is an input facsimile. The middle section is after one round of flipping pixels to optimize cluster representatives. The right section is after four rounds of flipping. Note the difference between large cluster text (letters i, n and o) and small cluster text (gh and th) in the right hand
 15 portion of Fig. 10.

[0087] Because the representative for a singleton (one-member) is identical in resolution to the scan, improvement is not attainable. Additionally, "furry" representation may take place where a vertical edge lies halfway between two verticals of the double-resolution pixel. One way to solve such problems is to define Bayesian
 20 prior probability distributions on representatives and incorporate them into the overall optimization using equation (5).

[0088] Chain codes may be used to determine the *a priori* distributions. A chain code is a string of letters N, E, S, W, for north, east, south, and west, representing the directions of boundary edges around a representative image. Edges
 25 are oriented so that black is on the left, meaning the edges are traversed counterclockwise around the outer boundary and clockwise around holes. Transition probabilities are compiled for all chain codes of length five, meaning the relative frequencies of the next letter after each possible string of length five. An exemplary method of using chain codes is described in U.S Patent 5,303,313 to Dance, which is
 30 incorporated herein by reference in its entirety.

[0089] Since boundary edges cannot double back on themselves, there are always three possible choices (straight, turn left or turn right) for each edge after the first. Hence, there is a total of $4 \times 3^5 = 972$ transition probabilities in the table.

Turing's rule is used for assigning probabilities to transitions that never occurred. That is, we assume that all non-occurring transitions had the same probability and that altogether they had the same total probability as the once occurring transitions.

[0090] The *a priori* probability $P[B^i]$ of a given representative image B^i is defined to be the product of the transition probabilities around all connected components of the boundary of the representative image. This *a priori* distribution penalizes furry representative images and rewards straight and smoothly curving representative images. The optimal representative image is thus defined to be the representative image B^i with maximum $P[C|B^i] \cdot P[B^i]$. The pixels on either side of an unlikely turn are marked as especially flippable, meaning that these pixels can be flipped even if the histogram argues against it.

[0091] The *a priori* probability $P[B^i]$ is much smaller than $P[C|B^i]$ for large clusters, and hence, rarely affects the choice of representative. For clusters with only two or three members, however, the *a priori* probability has an approximately equal voice in the outcome. For singletons, the *a priori* probability acts as a mild smoothing operation which improves straight strokes and staircasing along diagonals without rounding serifs.

[0092] Two different choices of training sets were used to compile the transition probabilities: the statistics from a clean postscript master, and (in a bootstrap approach) the statistics from the representatives for the large clusters (more than ten members) on the scanned document itself. No significant differences could be discerned between the two choices, even when the postscript master was the clean version of the scanned document.

[0093] C. Reclustering:

[0094] In exemplary embodiments, clusters are processed by decreasing order of their numbers of members. For each cluster, before representatives are computed, an attempt to merge the cluster with some larger cluster is performed. If cluster i is combined with some large cluster j , where large means more than three members, then the representative image B_j for the cluster j also serves as the representative image B_i for cluster i . If the larger cluster j is itself small, i.e., has no more than three members, however, then the combined cluster representative is redetermined using the members of both clusters. Alternatively, in various exemplary embodiments, merging cluster i with a larger cluster can be stopped when the size of the larger cluster gets down to

three. This alternative gives a significant increase in processing speed, sacrificing only a small amount of final image quality.

[0095] An exemplary reclustering is performed by using the connected component A_i denoting a single-resolution exemplar for cluster i , the representative image B_j denoting a double-resolution representative for cluster j , and the probability $P[A_i | B_j]$ as given by Eq. (3). In order to compare $P[A_i | B_j]$ against a preset threshold, $P[A_i | B_j]$ is normalized to account for the different sizes of connected components:

$$N[A_i | B_j] = (P[A_i | B_j])^{1/p},$$

where p is the number of pixels in the connected component A_i (aligned with the representative image B_j) that are within a sensor disk's radius of a black pixel in either the connected component A_i or the representative image B_j .

[0096] In one exemplary embodiment, a match occurs whenever $N[A_i | B_j]$ exceeds a threshold value. In various exemplary embodiments, the value threshold is 0.70. This threshold intuitively declares a match if the probability that the connected component A_i is a capture of the representative image B_j is at least the probability obtained if each pixel in the connected component A_i is predicted with probability 0.70. A slightly more aggressive threshold of 0.68 is used in the case that a cluster i is a singleton and a cluster j has at least four members. As a practical way to speed up the process, $N[A_i | B_j]$ is not determined if the bounding boxes for the connected component A_i and the representative image B_j differ too much in either width or height. If mergers are continued even when the larger cluster has fewer than four representatives, reclustering also saves some running time.

[0097] This reclustering improves output appearance significantly. For example, see Fig. 11 in which the left section is after four rounds of flipping including Bayesian priors. The center section is after reclustering. The right section is after breaking run-together letters. As can be seen from comparing the left and center sections, the m, M, c, te and th have improved after reclustering.

[0098] Reclustering can also improve compression performance by 10-30% on scanned and faxed documents, with a smaller percentage typical for flatbed scan and the larger percentage for typical 200dpi faxes. It has been found that super-resolution is important to reclustering performance. Thus, single-resolution representatives with single-resolution translations find only about 40% of the valid

mergers found by the double-resolution algorithm before starting to make mistakes. However, single resolution representatives with double-resolution translations find about 2/3rds of the valid mergers found by a fully double-resolution algorithm.

[0099] For fax inputs, there remain many singleton clusters, even after
 5 reclustering. Typically, half of these are run-together letters. In various exemplary
 embodiments of the invention, an additional step can be added to cope with this
 problem. In particular, for each singleton cluster, a final pass through its
 representative image is made, determining a sequence of "breakable positions" that
 attempt to break possible run-together letters. The "breakable positions" are as
 10 follows: a value of 2 is counted for each orthogonal adjacency, and a value of 1 is
 counted for each diagonal adjacency, between a column c and an adjacent column
 $c + 1$. The position between the column c and the adjacent column $c + 1$ is breakable
 if the total adjacency is no greater than 5, the last breakable position was at least 5
 columns to the left, and the total number of black pixels 2 and 3 columns to the left
 15 and right is sufficiently large (at least 6 on each of left and right). This check avoids
 breaking a horizontal line at every fifth column. The partial bitmaps are then matched
 (using the previous matching threshold as before) between successive breakable
 positions with previous clusters' representative images. If a successful match is found,
 then the partial bitmap is replaced by the representative of the larger cluster.
 20 Otherwise, the partial bitmap is passed along unchanged. Fig. 11 on the right section
 shows the result of the breaking step, in which of the run-together pairs gh, te and th,
 gh was successfully broken up and properly matched with letters g and h, respectively.

[00100] D. Assembling the Output:

[00101] The complete output image data is reassembled, replacing each
 25 connected component, or matched piece of a connected component, with its cluster's
 representative image. The position for the representative image is a most likely
 position found by first aligning the centers of bounding boxes of each connected
 component to be replaced and then testing the nine nearby double-resolution
 translations. All of these alignments may be redetermined, even though most of them
 30 were determined at an earlier step.

[00102] Various experimental results were conducted showing improvements
 in output image achieved by these various systems and methods. When compared to
 prior known techniques, error rates for a similar amount of clusters appears to be

lower. Of particular importance to overall quality and reproduction was use of super-resolved base line images and deskewing after super-resolution. To achieve reduced run time with only a slight reduction in image enhancement, it is possible to omit merging singletons with other singletons. Overall, it has been found that the systems and methods improve document images and a practical solution to high-quality scanning needs.

[00103] The foregoing description of the exemplary systems and methods for detection of this invention is illustrative, and variations in implementation will be apparent and predictable to persons skilled in the art. For example, while the systems and methods of this invention have been described with reference to desktop-captured images, any other type of image sensing device requiring accurate reconstruction of the underlying image can be used in conjunction with the systems and methods of this invention.

[00104] Thus, while the systems and methods of this invention have been described in conjunction with the specific embodiments outlined above, it is evident that many alternatives, modifications and variations will be apparent to those skilled in the art. Accordingly, the exemplary embodiments of the systems and methods of this invention, as set forth above, are intended to be illustrative, not limiting. Various changes may be made without departing from the spirit and scope of the invention.

[00105] For example, the methods and systems of this invention may also be useful for archival documents. If originals are no longer available, the methods and systems of this invention could improve the appearance of the existing image captures. Even if the originals are available, it may be more cost-effective to perform high-speed lower quality image captures and subsequently improve the image quality in software.